

MULTI-FACTOR APPLICATION SELECTION

The invention relates to application selection in a networked consumer environment. It finds particular application based on multiple factors associated with a desired content to be communicated with and/or rendered, a desired location for interacting with the content, and certain capabilities and conditions of applications and consumer electronic (CE) devices in the network and will be described with particular reference thereto. However, it is to be appreciated that the invention is also amenable to other applications.

Numerous advances have been made in CE devices, such as receiving, recording and content rendering devices for audio, video, and multimedia content. At the same time, communication technologies have also been advancing at a rapid rate. For example, digital streaming technology has provided users with the ability to obtain audio, video, and multimedia content via computer networks, such as the Internet. Significant advances are also occurring in other communication technology areas such as, e.g., home networking and automation, Internet access, and mobile wireless data services and devices.

Home environments contain a large variety of digital devices. For example, a home environment may include digital televisions (TVs), mobile phones, personal digital assistants (PDAs), satellite receivers, cable TV receivers, terrestrial antennas, DVD players, digital video recorders, etc. Furthermore, these devices may be interconnected in a "home network" by one or more home networking technologies.

The expression "home network" generally refers to a collection of interconnected apparatus in and around the home. An apparatus on the network can communicate with one or more of the other apparatus so as to provide distributed functionalities and synergy through interoperability, typically under control of software applications and macros. The network provides functionalities such as entertainment and education (e.g., audio and video play-out, electronic program guides (EPGs)), control (e.g., thermostat, lights, sprinkler, kitchen appliances) and monitoring (e.g., security system, baby monitor). A variety of software architectures may be implemented in a home network. Examples thereof include HAVi, Home API, UPnP, Jini, HomeRF, HomePNA, etc.

Several examples of a home network are discussed in PCT Published Patent Application WO 02/13463 A3. Devices within a locale, such as a home, are controlled by detecting the presence of an identified user within different areas of the locale. The devices within the areas are controlled in response to each identified user's preferences.

5 The locating and control devices may be stand-alone devices, or integrated within other electronic devices, such as TVs, stereos, computers, and so on. Also provided in this example of a home network are user task modules that suggest control actions based on the location of the user, the current context, and a profile of the user based upon the user's prior actions. The determination of each user's location is facilitated by the use of a tracker
10 module that the user carries about or other tracking devices. The user may determine the degree of automation to be applied at any time.

The rendering capabilities of different CE devices often differs. For instance, a digital TV may be capable of decoding high-resolution content streams, e.g., DVB MPEG-2 streams, while mobile phones may only be able to decode low-resolution
15 content streams, e.g., MPEG-4 movies transmitted via UMTS. In addition to heterogeneity in content rendering capabilities, the home networking capabilities of CE devices may also differ. For example, a TV may be equipped with a high bit rate IEEE 1394 interface, a PDA may be equipped with a moderate bit rate IEEE 802.11b wireless interface, and a mobile phone may be equipped with a low bit rate Bluetooth interface.

20 In this heterogeneous environment, a user should be able to access as much of the content available to the home on as many CE devices as possible. The content available from various service provider networks may include cable TV, satellite TV, wire and wireless telephone, and radio and TV broadcast. Content may also be available from some various storage devices in the home (e.g., DVD players, video and audio recorders,
25 CD players, a home archive (i.e., DVHS or other hard disk recorders, etc.)). The same content is often available in different formats or via different service provider networks.

Furthermore, the fact that CE devices in the home are interconnected enables a content provider to provide services to various devices, thus increasing the viewing time of the content provider's content. One way in which this can be achieved is
30 by offering content in different formats. The various formats may be transported via different networks. This enables the same content to be rendered on different CE devices even though each device requires a particular format, a particular network interface, or

both. Several examples of different formats over different networks include radio, TV, and the Internet. A content provider could, for example, broadcast a specific sports event via DVB-MPEG2, via UMTS-MPEG4, and via radio (e.g., FM/DAB). This enables a consumer to follow the event using his home cinema, mobile phone, or portable radio.

5 Multiple streams or files of data that, from a consumer's point of view, are related or very similar are referred to as -- "equivalent content" -- even though the streams or data are technically very different. Different types of content equivalencies may be distinguished as follows: 1) content that is exactly the same with different encoding characteristics (e.g., encoding scheme, bandwidth etc.); 2) content that is the same, but, for
10 example, the language or commentator are different; 3) the event from which the content originates is equal (e.g., a specific football match), but content is provided via different types of media (e.g., one is broadcast video media, while another is cable audio media). Note that these types of "equivalent content" streams are not required to arrive bundled in some access point in any way, nor are they required to come from the same service
15 provider.

Future CE networks in consumer environments will have a variety of devices and access to various content instances all with different quality and availability. This requires more sophisticated application selection mechanisms than are presently available to simplify the application selection process. This will allow users to concentrate
20 on selection of a desired content and a desired location in which to listen to and/or watch the content. Current application selection mechanisms only use content type-based application selection and then allow the user to influence which application is actually started. Other parameters associated with the selected content and/or the capabilities and conditions of the network are ignored.

25 For example, current desktop systems (e.g., Windows, KDE, Mac OS, etc.) determine applications that can be used for a given piece of content (i.e. normally a file) solely based on the type of the content. Windows uses the filename extension to determine the content type; whereas, KDE and MacOs take further information into account (e.g. the 'magic' starting bytes of the file). However, no other criteria are used.

30 With more and more types of CE devices available via networks and multiple users often sharing the same devices/resources, a sophisticated application selection mechanism is needed to select an application to communicate and/or render user-

selected content in a user-selected location, while also taking into account potential hindrances associated with available options. There is, therefore, a need for an apparatus and method for application selection based on a desired content and a desired location using multiple factors associated with the content and capabilities and conditions of the
5 network.

In one embodiment, an apparatus for selecting an application to deliver content in a networked consumer environment based on multiple factors is provided. The
10 apparatus includes a means for receiving an initial content selection, a means for i) accessing a list of registered applications for delivering content in the consumer environment and ii) accessing information identifying one or more types of content supported by each registered application in the list of registered applications, and a means
15 for i) identifying zero or more equivalent content references associated with the initial content selection, the combination of the initial content selection and the equivalent content references forming a list of selected content, ii) determining if any registered applications in the list of registered applications support any content types associated with the list of selected content, and iii) identifying each registered application in the list of
20 registered applications that supports any of the content types associated with the list of selected content as a compatible application.

In another embodiment, a method for selecting an application to deliver content in a networked consumer environment based on multiple factors is provided. An initial selection of content is received. Then, zero or more equivalent content references associated with the initial content selection are identified. The combination of the initial
25 content selection and the equivalent content references form a list of selected content. A list of registered applications for delivering content in the consumer environment is accessed. Information identifying one or more types of content supported by each registered application in the list of registered applications is accessed. Then, it is
30 determined if any registered applications in the list of registered applications support any content types associated with the list of selected content. Each registered application in the list of registered applications that supports any of the content types associated with the list of selected content is identified as a compatible application.

In still another embodiment, the method also includes receiving an initial location selection within the consumer environment to which the content is to be delivered. Information identifying the user interface requirements associated with each compatible application is accessed. A list of sink resources for delivering content in the consumer environment is also accessed. As well as one or more graphs for each compatible application. Each graph identifies sink resource requirements associated with the corresponding compatible application. It is determined if any sink resource in the initially selected location supports the user interface requirements and the sink resource requirements of any of the compatible applications. In this embodiment, each compatible application having sink resource requirements that are satisfied by one or more sink resources in the initially selected location is identified as a selected application.

In yet another embodiment, the method of the first embodiment also includes accessing a list of sources of content associated with the consumer environment. Each source that provides any of the selected content is identified. One or more graphs for each compatible application, each graph identifying source resource requirements associated with the corresponding compatible application are accessed. It is determined if any source associated with any of the selected content supports the source resource requirements of any of the compatible applications. In this embodiment, each compatible application having source resource requirements that are satisfied by any source associated with the selected content is identified as a selected application.

In yet still another embodiment, the method of the first embodiment also includes accessing information associated with user preferences for each compatible application. Information associated with previous executions of each compatible application is also accessed. As well as one or more graphs for each compatible application, each graph identifying resource requirements associated with the corresponding compatible application. A qualitative rating for each graph is determined based on at least one of the user preference information, the previous executions information, and information associated with a graph preference by the corresponding compatible application.

In another embodiment, the method of the first embodiment also includes accessing information associated with the allocated state of resources within the consumer environment. One or more graphs for each compatible application, each graph identifying

resource requirements associated with the corresponding compatible application are also accessed. It is determined if any available resource supports any of the resource requirements of any of the compatible applications. In this embodiment, each compatible application having resource requirements that are satisfied by one or more resources within
5 the consumer environment is identified as a selected application.

In still more embodiments, any combination of the previously described methods forms an additional embodiment.

One advantage of the invention is a content instance having lowest cost can be selected from equivalent contents and rendered to a given content rendering device in a
10 networked consumer environment.

Still another advantage is, in multi-device clusters, the best possible sink will be chosen/proposed automatically; taking into account how much screen space is actually available.

Yet another advantage is an application is selected that is best capable of
15 taking advantage of special features of a source (or can best handle its disfeatures).

Still yet another advantage is application selection is adapted to typical user selection over time.

Another advantage is no application is started that would over exhaust the available resources, even if that application could in principle provide better quality than
20 the application selected in its stead.

Other advantages will become apparent to those of ordinary skill in the art upon reading and understanding the following detailed description.

25 The drawings are for purposes of illustrating exemplary embodiments of the invention and are not to be construed as limiting the invention to such embodiments. It is understood that the invention may take form in various components and arrangement of components and in various steps and arrangement of steps beyond those provided in the drawings and associated description. Within the drawings, like reference numerals denote
30 like elements.

FIG. 1 is a block diagram of an embodiment of a communication environment incorporating the invention.

FIGs. 2A-2E are flowcharts showing various embodiments of a method for application selection in a consumer environment network using multiple factors associated with desired content, desired location, and capabilities and conditions of the network.

FIG. 3 is a flowchart showing another embodiment of a method for application selection using multiple factors.

FIG. 4 is a realized graph of end-to-end communication in the consumer environment when audio-visual (AV) content is desired after the application selected to provide the content is running.

FIG. 5 is a realized graph of end-to-end communication in the consumer environment when a video conference is desired after the application selected to conduct the conference is running.

With reference to FIG. 1, a communication environment 10 includes one or more consumer environments 12, one or more content providers 14, and one or more content/service providers 16. The communication environment 10 encompasses infrastructures for wired and wireless telephone communication, satellite communication, cable TV communication, terrestrial radio and TV broadcasts, and other communication networks capable of distributing content to consumers and communication between consumers. Typically, a consumer environment 12 is a home or apartment, a vehicle, an office, or a business. An example of a business environment that may implement this invention is a business that sells audio, video, or multimedia devices to consumers. Another example is a business that makes audio, video, or multimedia content available to consumers via a network of such audio, video, or multimedia devices (e.g., a library or an Internet cafe).

The content/service providers 16 include broadcasters, Internet service providers (ISPs), local and long distance telephone companies, wireless telephone service providers, cable TV service providers, satellite TV service providers, and other types of service providers that offer audio, video, or multimedia content to consumers. Content/service providers 16 provide access to content from multiple content providers 18 and a path for communication between multiple consumer environments 12. Content/service providers 16 may also provide content like a content provider 14. Many

content/service providers 16 offer subscription services to consumers. Alternatively, or in addition, content/service providers 16 may offer such specialty services as pay-for-access content, on-demand content, and downloads (e.g., movies, music, e-books). There are various business models for the various types of content/service providers 16 and
5 competition often leads to various pricing schemes for the same type of service.

A given content provider 14 may provide equivalent content to one or more content/service providers 16 via the Internet, a wired interface, or a wireless interface. The interface between a content provider 14 and a content/service provider 16 can implement any communication protocol capable of transporting the desired content. The content
10 provider 14 may communicate the same content via multiple types of communication protocols. Similarly, a given content/service provider 16 may provide equivalent content to the consumer environments 12 via the Internet, a wired interface, or a wireless interface. The interface between a content/service provider 16 and a consumer environment 12 can implement any communication protocol capable of transporting the desired content. The
15 content/service provider 16 may communicate the same content via multiple types of communication protocols. In general terms, the transport methods for communication between consumer environments 12, content providers 14, and content/service providers 16 include cable, satellite, terrestrial, and wired or wireless telephone. This includes RF broadcast, satellite communication, cellular communication, wired, and fiber optic
20 technologies.

The consumer environment 12 includes a network 18 and one or more clusters 20. A cluster 20 is a group of consumer electronic (CE) devices associated with a particular location (e.g., a room or a portion of a room) within the consumer environment 12. Certain CE devices within the one or more clusters 20 are in communication with each
25 other via the network 18. The network 18 may use any wired or wireless communication protocol (e.g., IEEE 1394, Ethernet, IEEE 802.11b, Bluetooth, etc.) capable of transporting content compatible with the surrounding networked devices. The network 18 may also be a hybrid combination of networks utilizing a combination of protocols and a combination of wired and wireless technologies (i.e., wired, fiber optic, low-level RF, IR).

30 Clusters 20 can include any type of CE device and can be arranged in various configurations. For purposes of this invention, a CE device is viewed in terms of its functionality. For example, a given CE device may include external communication,

processing, source, or sink functions or any combination of these functions. A CE device is viewed as a resource for any and all functions that it provides to the consumer environment 12. Therefore, this invention deals with external communication resources, processing resources, source resources, and sink resources. Any type of CE device that provides access to a content/service provider 16 via any infrastructure in the communication environment 10 (e.g., receivers, tuners, antennas, modems) contains an external communication resource. Any type of CE device that processes audio, video, or multimedia content (e.g., set-top boxes, decoders, transcoders, amplifiers, image processors, computers) contains a processing resource. Any type of CE device that stores or captures audio, video, or multimedia content (e.g., CD or DVD players, audio or video cassette players, disk drives, cameras, microphones, 2-way radios, telephones) contains a source resource. Any type of CE device that can present or render audio, video, or multimedia content to a consumer (e.g., displays, projectors, 3D projectors, speakers, radios, telephones, televisions (TVs), personal digital assistants (PDAs)) contains a sink resource.

Three examples of cluster configurations with various combinations of resources are shown. A first cluster 20a includes one or more sink resources 22, one or more source resources 24, and a processing resource 26, each in communication with the network 18. The first cluster 20a also includes one or more sink resources 22, one or more source resources 24, and one or more external communication resources 28, each in communication with the processing resource 22. Each external communication resource 28 is also in communication with one or more content/service providers 16. In one respect, within the context of a given consumer environment 12, an external communication resource 28 acts a source resource for content provided to the consumer environment 12 by content/service providers 16.

As discussed above, a CE device can include any combination of one or more resources. For example, in the first cluster 20a, the resources in communication with the processing resource 26 could be combined in a CE device to provide functionality that is available in certain configurations of computer systems. This computer system configuration may be characterized as a computer chassis (i.e., processing resource), a monitor and speakers (i.e., sink resources), a camera and a microphone (i.e., source resources) and a modem (i.e., external communication resource).

A second cluster 20b includes one or more sink resources 22 and one or more source resources 24, each in communication with the network 18. A third cluster 20c includes one or more sink resources 22, each in communication with the network 18. Note that the second and third clusters 20b, 20c do not include processing resources 26 or external communication resources 28. Therefore, these clusters rely on one or more centralized processing resources 30 and one or more centralized external communication resources 32. The centralized processing resources 30 and centralized external communication resources 32 are shared resources that are not associated with any particular cluster. The centralized processing resources 30 and centralized external communication resources 32 provide the same functionality described above for the processing resource 26 and external communication resources 28 in the first cluster 20a. Alternatively, the network 18 can be configured so that a processing resource 26 and/or an external communication resource 28 associated with a cluster is shared with one or more sink resources 22 and/or one or more source resources 24 in other clusters.

The consumer environment 12 may also include one or more mobile devices 34 (e.g., portable radios, portable CD or DVD players, portable TVs, PDAs and other types of portable computers). Mobile devices 34 are CE devices that include various combinations of communication, processing, source, and/or sink resources or functionality and are portable. As such, any mobile device provides standalone functionality to users and may be used outside the local consumer environment. For example, a portable TV can be used in an external environment, such as a vehicle or a park. Similarly, a PDA can be used in an external environment, such as an office. Each mobile device 34 is also a network resource within its local consumer environment and includes a local wireless interface (e.g., low-level RF, IR) to communicate with the network 18. When a mobile device 34 is within communication range, it provides its resources or functionality to the network 18.

At least one networked resource in the consumer environment 12 includes an application manager that provides multi-factor application selection. Preferably, each sink resource includes an application manager managing all applications that present content on that sink. A user interface function is associated with the application manager. This user interface function allows users to select desired content for communication and/or rendering and a desired location for communicating and/or rendering the content.

Any user interface associated with multi-factor application selection includes a display (e.g., monitor, character readout, indicators) and an input component (e.g., keyboard, mouse or other type of pointing device, keypad or other type of switches/controls). Different devices can provide the user interface using different types of components and in
5 different levels of sophistication. Typically, the display and input components are required by communication, processing, source, and/or sink resources of networked devices and re-used in conjunction with the user interface function associated with multi-factor application selection.

If the application manager is only provided at one resource in the consumer
10 environment 12, a centralized control strategy is implemented for selection of applications. If the application manager is provided on multiple resources in the consumer environment 12, either a master-slave control strategy or a synchronized control strategy may be implemented. Other control strategies are also possible. In any control strategy, the user interface function may be implemented using local and remote user interfaces. In other
15 words, the application manager can be controlled using the local user interface of the resource on which the application manager resides or it can be controlled using a remote user interface of another networked resource within the consumer environment 12. For the master-slave control strategy, any resource running the application manager can be designated master and, if necessary, the master status can be transferred. For the
20 synchronized control strategy, information regarding content communication and rendering and the current status of various networked resources is communicated to each application manager to coordinate application selection across the consumer environment 12.

The application manager selects an appropriate application for a content the user wants to communicate with and/or have presented. The application manager takes
25 several factors into account to choose the "optimal" application to communicate and/or present the content to the user in the specific situation. These factors include: 1) content type, 2) capabilities of the intended sink, 3) availability of equivalent versions of the content and capabilities of the sources providing the equivalent versions, 4) user preferences from previous application selections, and 5) allocation and availability of
30 resources and bandwidth. Applications will express their streaming requirements to the application manager in advance as a graph. A graph is a collection of (abstract) resource nodes (i.e., resource requirements) connected via edges. Graphs can have tags assigned

that further specify the required functionality. A graph is realizable if an actual resource (i.e., functional resource in a specific device) having the required functionality can be assigned to each of the resource nodes and the required bandwidth between the nodes can be allocated. Therefore, a realizable graph is a graph with all resources available.

5 Quantitative values for selection criterion may be used to pre-select the realizable graph with the highest overall rating. Multiple factors are used to ultimately determine which application to start.

The application selection process begins, for example, when a user in the consumer environment uses the user interface of a networked resource to express his or her desire to be presented some kind of content (e.g., he wants to watch a movie or to listen to a football match broadcast). The consumer environment must then find an application which can present the content to the user with the maximal possible quality. Since quality criteria might be different between users and user intentions differ from one moment to another, typically the application manager incorporates preferences that can be adjusted by users and/or allows user interaction during the application selection process.

An example of the application selection process based on available resources is where the user selects video content for rendering. The video content is available as an MPEG2 stream from the Internet or from a DVD so the content type is the same. The application to show the stream from the DVD is preferably the application that came with the DVD player (or any other application which can make use of the advanced features of a DVD player like switching different language soundtracks). The application to show the content from the Internet might contain special measures to improve picture quality, which normally suffers due to the higher compression in streams from the Internet. Depending on which source the application manager finally chooses, one of the applications is ultimately selected.

Another example of the application selection based on available resources is where the user selects a videoconference for communicating content to another consumer environment. A videoconference application might require at least 320 x 240 pixels of screen space to present all of the UI elements (e.g., local and remote camera views, control buttons) associated with the application. If the space required for all the UI elements is not available, at least the video for the remote camera view must be displayable and the UI must support some kind of tabbing widget.

Still another example of the application selection based on available resources is where the user selects another user as the content he or she wants to get presented (communication between the users). If there is no resource in the consumer environment with space available for any video from the cameras, the application manager should not suggest a videoconference. Rather, the application manager should select an audio conference or some type of text messaging.

An example of the application selection process based on the allocation status of available resources is where a son of the family selects a movie for rendering and one source of the movie requires a resource that is currently in use by another application. The highest quality version of the movie the son wants to watch is available from a DVD in the home storage, a lower quality one from the Internet. However, the DVD version of the movie is only available in a stream format that is not understood by the display where the son sits. Therefore, the DVD application requires processing of the stream by a transcoder. A transcoder is a device to convert a stream (i.e., audio and/or video) from one, for example, format, protocol, size, and/or resolution to another. However, all transcoders are currently in use. The application manager selects the Internet application and makes that version available to the son.

An example of the application selection process based on user preferences is where the parents prefer a well-known stable version of the DVD application, while the son normally prefers the newest version. As shown by the examples, the final selection of content, source resources, sink resources and application is mutually dependent.

With reference to FIGs. 2A-2E five embodiments of a method for multi-factor application selection are provided. FIG. 2A provides an embodiment of the method that includes an equivalent content service with content type-based application selection. FIGs. 2A and 2B provide an embodiment of the multi-factor application selection method that combines content type-based application selection and sink-based application selection. FIGs. 2A and 2C provide an embodiment that combines content type-based application selection and source-based application selection while also including equivalent content and source finder services. FIGs. 2A and 2D provide another embodiment of the method. This embodiment combines content type-based application selection and preference-based application selection. In still another embodiment, FIGs.

2A and 2E provide a method that combines content type-based application selection and availability-based application selection.

In preparation for multi-factor application selection, each application associated with content communication and/or content rendering registers with the system.

5 This forms a set of registered applications. As part of the registration process, each application specifies its required resources, UI requirements, and content types that it supports. This information is also consolidated in the form of resource graphs that reflect various paths for end-to-end communication that are supported by the registered application. Initially, these graphs contain placeholders or nodes, for example, for sink
10 resources and locations and for content and content sources (i.e. the initial attributes at the nodes reflect the capabilities of the registered application). In certain embodiments, tracking processes are also associated with the application manager to maintain records of user preferences, previous uses of the registered applications, and current status and allocation of networked resources.

15 With further reference to FIG. 2A, one embodiment of a method for multi-factor application selection starts at a step 102 when a user initially selects content. Typically, the user will use an already running system application to specify this request. This "initiating application" will contact the application manager and pass it the requested content. Next, at a step 104, the equivalent content finder checks for equivalent content.
20 Calculation and provisioning of content equivalencies may be retrieved from some third party service. This may result in a list of content that is considered equivalent to the originally requested content (e.g. radio broadcast of a football game as alternative to a TV broadcast). A qualitative rating between 0 and 100% (i.e., 0.00 to 1.00) may be determined for each returned equivalent content as well as the initially selected content.
25 The content with the highest quality establishes the baseline (i.e., 100% or 1.00) for the qualitative rating. The application manager may automatically limit further processing to the initially selected content and equivalent content that exceeds a certain threshold qualitative rating. Additionally, or alternatively, the application manager may allow user selection of one or more items for further processing from a list of the initially selected
30 content and the returned equivalent content.

The content selected for further processing in the step 104 and the set of registered applications (step 106), along with the content types supported by each

registered application (step 108) are provided for content type-based application selection process in a step 110. Registered applications that do not support the content types associated with the returned equivalent content are filtered out in the step 110. The content type-based application selection process typically comes up with several possible combinations of equivalent content and registered applications and identifies one or more compatible applications (step 112). To avoid presenting a long list of compatible applications to the user and also to facilitate some automatic tasks (e.g. improving the application performance if new equivalent content or resources become available) a qualitative rating that indicates how much the user might appreciate the compatible applications may be associated with the result. The application manager may automatically select the compatible application with the highest qualitative rating. Alternatively, the application manager may allow user selection of an application from a list of the compatible applications. The application manager may automatically limit the list of compatible applications to applications that exceed a certain threshold qualitative rating.

As described above, information associated with each registered application also identifies UI requirements (step 114) and graphs (step 116) for the application. This information is available to the application manager and may be introduced in other embodiments of the method for multi-factor application selection. For these additional embodiments, the application manager includes a graph mapper process. The results of the equivalent content finder (step 104) and the results of the content type-based application selection process (step 110) may be provided to the graph mapper process to fill in the nodes of graphs for the compatible applications.

With further reference to FIGs. 2A and 2B, another embodiment of a method for multi-factor application selection starts at the step 102 and continues through FIG. 2A as described above. In this embodiment, the compatible applications identified in the step 112 are carried forward for further processing based on sink resources. At a step 118, the user initially selects a location in the consumer environment in which he or she wants to interact with the content. Typically, the initial location selection occurs along with the initial content selection (step 102) through the "initiating application." Like the initial content selection, the initial location selection is passed to the application manager by the "initiating application." The initial location is typically specified as a cluster of

devices (e.g., the TV/high fidelity corner in the living room). Alternatively, the initial location can be a specific device (e.g., PDA). The compatible applications (step 112), UI requirements associated with the compatible applications (step 114), graphs associated with the compatible applications (step 116), initial location selection (step 118), and sink resources (step 120) are provided for sink-based application selection in a step 122.

If the initial location was a cluster of devices and not an individual sink resource, the application manager queries a resource registry for the cluster to retrieve the individual sink resources in the cluster. The capabilities of the sink resources are compared to the UI requirements of the compatible applications. UI requirements, for example, contain information about required screen space, widget types, and audio creation capabilities. Typically, applications require some way to display a UI to the user (i.e., displaying content is also some kind of UI). Applications should be programmed more or less independent of the type of resource used for the UI, but most applications will nevertheless have minimum requirements for the UI used to present or render the selected content. The application manager compares the UI capabilities of the sink resources to the UI requirements of the compatible applications.

Applications that are not compatible with sink resources in the selected location are filtered out in the step 122. The sink-based application selection process typically comes up with several possible combinations of selected content, registered applications, and sink resources and identifies one or more selected applications (step 124). If qualitative ratings are used in the application selection process, the application manager may automatically select the application with the highest qualitative rating. Alternatively, the application manager may allow user selection of an application from a list of the selected applications. The application manager may automatically limit the list of selected applications to applications that exceed a certain threshold qualitative rating. At this point, the results of the sink-based application selection process (step 122) may be provided to the graph mapper process to fill in nodes of the graphs for the selected applications.

With further reference to FIGs. 2A and 2C, another embodiment of a method for multi-factor application selection starts at the step 102 and continues through FIG. 2A as described above. In this embodiment, the compatible applications identified in the step 112 are carried forward for further processing based on the source of selected content. Note that the initial content selection (step 102) and the equivalent content finder

(step 104) described above for FIG. 2A also contribute to source-based application selection. The content selected for further processing in the step 104 and available sources of content (step 126) are provided to a content source finder in a step 128. Sources include source resources, addresses or identifiers associated with other consumer environments, and channels, sub-channels, addresses, or other types of identifies associated with content providers. Available sources include source resources currently on the network in the consumer environment and other consumer environments and content providers that are accessible via current subscriptions with service providers.

Applications may be source specific (e.g. prefer to work with DVD-players). Therefore some applications may restrict sources in their graph. Therefore the application manager must test for each application graph, if there exists a source of the required type/types providing the selected content. Applications may prefer specific kinds of sources even though the content types the application handles can come from other sources as well. An example is a special DVD-watching application that can use the DVD-player capabilities to switch languages or to show the bonus' tracks. A standard video watching application is probably not able to do this, but may be better for showing non-DVD video. The content source finder finds all sources that can provide the content selected for further processing. The content source finder returns a list of source s associated with the selected content for further processing.

The compatible applications (step 112), graphs associated with the compatible applications (step 116), and the resulting list of sources from the step 128 are provided for source-based application selection in a step 130. Applications that specify a required source that does not provide the selected content are filtered out in the step 130. The source-based application selection process typically comes up with several possible combinations of selected content, registered applications, and sources of the selected content and identifies one or more selected applications (step 124).

If qualitative ratings are used in the application selection process, the application manager may automatically select the application with the highest qualitative rating. Alternatively, the application manager may allow user selection of an application from a list of the selected applications. The application manager may automatically limit the list of selected applications to applications that exceed a certain threshold qualitative rating. At this point, the results from the equivalent content finder (step 104), the results

from the content source finder (step 128), and the results of the source-based application selection process (step 130) may be provided to the graph mapper process to fill in nodes of the graphs for the selected applications.

With further reference to FIGs. 2A and 2D, another embodiment of a method for multi-factor application selection starts at the step 102 and continues through FIG. 2A as described above. In this embodiment, the compatible applications identified in the step 112 are carried forward for further processing based on preferences. The compatible applications (step 112), graphs associated with the compatible applications (step 116), user preferences associated with the compatible applications (step 131), and previous uses or executions of the compatible applications (step 132) are provided for preference-based application selection in a step 134. In addition, the graph mapper process, through feedback from the preference-based application selection process (step 134), may fill in qualitative ratings in the graphs (step 116) so that the graphs may be ranked according to the user preferences.

The system itself maintains a database of usage preference for each application relative to content type and user. Registered applications that are preferred by the user for specific kinds of content are given a higher qualitative rating (e.g., typically based on usage frequency). If equivalent content from different sources varies in quality (e.g. broadband MPEG2 versus highly compressed MPEG4) this is also reflected in the qualitative ratings. Even if several applications are capable of presenting the same content, the user may prefer a specific application (e.g. because this application is the only one capable of editing the content).

Applications may specify preference percentages with graphs (i.e. qualitatively ranking its graphs based on which graph the application likes best) when registering with the system. These qualitative ratings may be taken into account by the preference-based application selection process to filter out graphs that are rated below a certain threshold qualitative rating when deciding which options to propose to the user.

The preference-based selection process 134 includes consideration of one or more of the following qualitative ratings: i) a content qualitative rating, ii) a UI qualitative rating, iii) an application qualitative rating, iv) a user qualitative rating, v) a graph qualitative rating, and vi) a graph mapper qualitative rating. The content qualitative rating is a number assigned by the equivalent content finder. The content qualitative rating

represents the overall quality of the content (e.g., consideration of size, resolution, frame rate, etc.). The UI qualitative rating is a number assigned by a UI-shell associated with the application manager. The UI qualitative rating expresses which sink devices the UI shell would prefer based on, for example, remaining available screen space. The application qualitative rating is a number gathered over time by the application manager. The application qualitative rating expresses how often the user chose a particular application for this specific content type on previous occasions. The user qualitative rating: is a number specified explicitly by the user to express his or her preferences for a specific application (e.g., the cutting edge version vs. the trusted old version). The graph qualitative rating: is a number representing how appropriate a particular graph is for the application. Typically, the graph qualitative rating is specified by the application designer to give preference to "normal" graphs over "resources are low" graphs. The graph mapper qualitative rating represents how "easy" it would be to allocate/assign all the resources requested by the graph.

All of these qualitative ratings may be implemented or any combination thereof. Each qualitative rating that is implemented is used to determine an aggregate qualitative rating for each potential application. Typically, the implemented qualitative ratings are multiplied together for a given application to determine the aggregate qualitative rating of that application.

The compatible applications are ranked according to preferences in the step 122 and one or more selected applications are identified in a step 124. The application manager may automatically select the application with the highest qualitative rating. Alternatively, the application manager may allow user selection of an application from a list of the selected applications. The application manager may automatically limit the list of selected applications to applications that exceed a certain threshold qualitative rating. At this point, the results of the preference-based application selection process (step 134) may be provided to the graph mapper process to fill in nodes of the graphs for the selected applications.

With further reference to FIGs. 2A and 2E, another embodiment of a method for multi-factor application selection starts at the step 102 and continues through FIG. 2A as described above. In this embodiment, the compatible applications identified in the step 112 are carried forward for further processing based on availability of resources in

the consumer environment. The compatible applications (step 112), graphs associated with the compatible applications (step 116), and a current state and allocation of resources (step 136) are provided for availability-based application selection in a step 138.

Each registered application provided all its graphs to the application manager when the application was registered in the system. Now, the application manager tests these graphs for realizability (i.e., availability of resources) to find out which application can be used to present the content to the user. Testing for realizability is time consuming. Therefore, it is desirable to test as few graphs as possible. The application manager process, through feedback from equivalent content finder (step 104), the content type-based selection process (step 110), and the availability-based application selection process (step 138), and limits the graphs considered to those that are compatible with resources that are currently available for allocation. Therefore, based on the information selected so far, graphs are filtered out. For example, if no sink resource is usable for a particular graph, the graph is filtered out.

A complete description of the required resources (including streaming interconnections) is tested for availability given the current allocation status of the available resources. The execution of an application requires bandwidth and may require a specific resource (e.g. a tuner or a transcoder). This information is read from nodes of the application graphs and, if allocation of the required bandwidth or the required resources fails, the graph is filtered out. For example, if no resource is available for a particular graph or if the required space or capacity is not available in any resource for a particular graph, the graph is filtered out. If all of the graphs for a particular application are filtered out, the application cannot be started and it is also filtered out.

The availability of the (streaming) resources required for each graph is tested by the graph mapper process. The graph mapper receives the application graphs tagged with content and possible sinks and tests the availability of the sources for the content. Also, depending on the available and requested features of the content (i.e., size, frame rate, etc.), the graph mapper may need to include transcoders or other processing resources in the graph. The graph mapper may return a availability-quality rating which is a value indicating how much 'effort' this realization was (e.g. how many transcoders were introduced, is the network bandwidth becoming scarce).

For example, two graphs may exist for a given application, one graph using a satellite tuner as a source resource, the other using a cable tuner. If both tuners can serve two streams each and the satellite tuner already serves one stream, the graph using the cable tuner is given a higher rating assuming everything else is equal. Similarly, if the availability of a particular resource required by an application needs additional resources (e.g. transcoders) inserted or if it uses up a scarce resource (e.g. much bandwidth)) the application may be assigned a low qualitative rating. While, if another sink resource is already compatible with the content format delivered by an application, its allocation status will return a high qualitative rating.

The aggregation of various quality ratings results in one quality rating per graph and thereby also one for the usefulness of each application in the current situation. The quality ratings for applications are used to order possible application-content-location combinations to the user in an arrangement that emphasizes applications expected to be the most useful. But also the quality ratings for graphs can serve a useful purpose. Testing the availability of all required resources in a graph is a time consuming process. Also typically many possible graphs will exist for a content-location combination selected by the user. The graphs might be tested in a specific order by the graph mapper (e.g. based on the preference percentages returned by the previous steps). Testing of additional graphs may be stopped after a certain amount of time has elapsed or when the qualitative ratings for the remaining graphs are below a certain qualitative threshold. This is to ensure timely startup of an application, even if many alternative applications/graphs exist.

Applications that specify a required resource that is not available are filtered out in the step 138. The availability-based application selection process typically comes up with several possible combinations of selected content, registered applications, and available resources and identifies one or more selected applications (step 124). If qualitative ratings are used in the application selection process, the application manager may automatically select the application with the highest qualitative rating. Alternatively, the application manager may allow user selection of an application from a list of the selected applications. The application manager may automatically limit the list of selected applications to applications that exceed a certain threshold qualitative rating.

With reference to FIG. 3, another embodiment of a method for multi-factor application selection consolidates the several embodiments described above and shown in

FIGs. 2A-2E. In the embodiment being described, a pipeline from the registered applications (step 106) to the one or more selected applications (step 124) is shown without repeating the detailed steps associated with each of the previously described embodiments. In the step 110, initial content selection, equivalent content, and content type are factors
5 that contribute to content type-based application selection. Applications that do not support the content types associated with the selected content are filtered out in the step 110. Compatible applications are carried forward for further processing.

In the step 122, UI requirements associated with the compatible applications, graphs associated with the compatible applications, initial location selection,
10 and sink resources are factors that contribute to sink-based application selection. Applications that do not support sink resources in the selected location are filtered out in the step 122. Compatible applications are carried forward for further processing.

In the step 130, graphs associated with the compatible applications and sources for the selected content are factors that contribute to source-based application
15 selection. Applications that require a source that cannot provide the selected content are filtered out in the step 130. Compatible applications are carried forward for further processing.

In the step 134, graphs associated with the compatible applications, user preferences associated with the compatible applications, and previous uses or executions of
20 the compatible applications are factors that contribute to preference based application selection. Applications are ranked according to preferences and low-ranked applications may be filtered out in the step 134. Compatible applications not filtered out are carried forward for further processing.

In the step 138, graphs associated with the compatible applications and a
25 current state and allocation of resources within the consumer environment are factors that contribute to availability-based application selection. Application that require resources that are not available are filtered out in the step 138. One or more selected applications are ultimately identified in the step 124. The application manager may automatically select the application with the highest qualitative rating. Alternatively, the application manager may
30 allow user selection of an application from a list of the selected applications. The application manager may ~~automatically limit~~ the list of selected applications to applications that exceed a certain threshold qualitative rating.

For the embodiment being described, steps 110, 122, 130, 134, and 138 can be performed in any order and in any combination. However, each arrangement has its advantages and disadvantages. For example, having the user select early from the available applications the one fitting his intentions would reduce the number of potential applications early on and therefore reduce the amount of necessary tests and speed up application startup. On the other hand, one wants to avoid proposing applications to the user that cannot be run if selected. Performing availability selection very late reduces the number of graphs to be tested and may even allow applications to influence their graphs before testing but restricts the degrees of freedom for the graph mapper process and might require testing graphs for applications that are not ultimately started.

When the user selects a content the task of the system is to find the set of applications which can be used for presenting the content in the current situation to the user with an ordering based on overall estimated quality of the resulting presentation. The selection process described above typically comes up with several combinations of selected content, sources for the selected content, sink resources in the selected location, and compatible applications. To avoid presenting a long list of compatible applications to the user and also to facilitate some automatic tasks (e.g. improving the application performance if new equivalent content or resources become available), qualitative ratings are introduced in one or more stages of the selection process. For example, the equivalent content finder assigns a percentage from 0 to 100 to all equivalent content it returns. The number associated with selected content represents some kind of quality rating for the equivalent content with respect to the initially selected content. The estimated preference of a user to select a given application is also returned as a value that indicates how probable the user will want to use the application with the selected content types. This value is typically adapted over time from user behavior.

As described above, the graph mapper process return a percentage between 0 and 100 to indicate how well it could get all required resources. The application manager process will also integrate the preference ratings for graphs specified by the applications during the registration process. Finally, all these qualitative ratings are multiplied together. Using the aggregated rating, the selected applications are ranked and either the highest rated application is started automatically or the user is given the choice out of most highest rated content-source-sink-application combinations.

This choice of the user will be enabled by the "initiating application" which receives the ordered (and possibly truncated) list of possible application-content-location triple combinations from the application manager. After the user has made this choice via the user interface of the "initiating application" (or if that application has decided on its own) it will name the selected triple combination to the application manager, which will then take care of starting the appropriate application as ordered. However, non-interactive initiating applications (e.g., timer-based auto starter) can circumvent this external selection by ordering the application manager to select among the possible applications on its own instead of passing an ordered list to the "initiating application." The application manager will then typically choose the application-content-location combination with the highest aggregated quality rating.

With reference to FIG. 4, an end-to-end communication path 150 shows a realized graph for viewing audio-video content, such as a movie. This is a simplified version of the realized graph for the application selected to deliver the content. The actual graph includes specific device identifiers (e.g., device type, address, location), specific source identifiers (e.g., channel, sub-channel, frequency), and performance parameters (e.g., resolution, screen size, frame rate, bandwidth). As shown, the source of the audio-video content is a content provider 14 via a content/service provider 16. An external communication resource 32 in the consumer environment 12 receives the audio-video content. A processing resource 30 conditions the audio-video stream for two sink resources 22 (i.e., a video display and audio speakers) located in a particular cluster 20.

With reference to FIG. 5, another example of an end-to-end communication path 160 in the consumer environment 12 shows a simplified realized graph for a video conference. Here, the source of audio-video content is another consumer environment 12 via a content/service provider 16. An external communication resource 32 in the consumer environment 12 receives the audio-video content and provides it to a particular cluster 20. A processing resource 26 in the cluster 20 conditions audio-video streams from the external consumer environment 12 for two sink resources 22 (i.e., a video display and audio speakers). For example, the processing device 26 may be responsible for converting audio-video streams from/to H.323. Two source resources 24 (i.e., a camera and a microphone) provide audio-video content to the processing resource 26. The processing resource 26 conditions the audio-video stream from the source resources 24 and audio-

video content is communicated back to the external consumer environment 12 via the same path as that for receiving content.

5 While the invention is described herein in conjunction with exemplary embodiments, it is evident that many alternatives, modifications, and variations will be apparent to those skilled in the art. Accordingly, the embodiments of the invention in the preceding description are intended to be illustrative, rather than limiting, of the spirit and scope of the invention. More specifically, it is intended that the invention embrace all alternatives, modifications, and variations of the exemplary embodiments described herein that fall within the spirit and scope of the appended claims or the equivalents thereof.